

EXHIBIT C

SUSMAN GODFREY L.L.P.

A REGISTERED LIMITED LIABILITY PARTNERSHIP
ONE MANHATTAN WEST
NEW YORK, NEW YORK 10001-8602
(212) 336-8330
FAX (212) 336-8340
WWW.SUSMANGODFREY.COM

SUITE 5100
1000 LOUISIANA STREET
HOUSTON, TEXAS 77002-5096
(713) 651-9366

SUITE 1400
1900 AVENUE OF THE STARS
LOS ANGELES, CALIFORNIA 90067-6029
(310) 789-3100

SUITE 3000
401 UNION STREET
SEATTLE, WASHINGTON 98101-2683
(206) 516-3880

ZACH SAVAGE
DIRECT DIAL (212) 729-2022

E-MAIL ZSAVAGE@SUSMANGODFREY.COM

July 22, 2024

VIA ELECTRONIC MAIL

Jared B. Briant
FAEGRE DRINKER BIDDLE & REATH LLP
1144 15th Street, Suite 3400
Denver, Colorado 80202
(303) 607-3588
jared.briant@faegredrinker.com

Annette L. Hurst
ORRICK, HERRINGTON & SUTCLIFFE LLP
The Orrick Building
405 Howard Street
San Francisco, CA 94105
(415)773-5700
ahurst@orrick.com

Re: *The New York Times Company v. Microsoft Corporation, et al.*,
Case No.: 23-cv-11195-SHS

Counsel:

I write regarding Microsoft's Responses and Objections to The Times's First Set of Interrogatories and Second Set of Requests for Production, served on June 21, 2024 (the "R&Os"). Please agree in writing to remedy the deficiencies identified below or offer a time to meet and confer by no later than July 26.

1. Responses and Objections to The Times's First Interrogatories

a. Interrogatories for which Microsoft's responses are incomplete (Interrogatories Nos. 1, 3, 7, 8, 9)

Interrogatory 1 requests identification of individuals with knowledge of (i) Defendants' Training Datasets, including the existence of Times Content and other copyrighted content therein, (ii) all phases of training of the Text Generation AI Models, (iii) the business practices regarding the Text Generation AI Models, and (iv) the commercialization and use of the Text Generation AI Models. Microsoft limits its response to Interrogatory 1 to "individuals who are knowledgeable regarding the timeframes, contracts, investments, and/or computing infrastructure used for development with respect to OpenAI." R&Os at 6. There are several problems with this response. First, Microsoft's response fails to identify individuals with knowledge of sub-parts (i) Defendants' Training Datasets and (ii) all phases of training. The Times understands Microsoft's position to be that Microsoft did not "pre-train" Defendants' AI models, but that Microsoft has "post-trained" those models. Nevertheless, Microsoft must still identify those individuals within Microsoft who do have knowledge of all the training datasets and phases of training. Are Kate Cook and Mikhail Parakhin also the individuals at Microsoft who are most knowledgeable regarding sub-parts (i) and (ii)? If so, please amend your response accordingly. If not, please identify which individuals possess such knowledge. Second, as for sub-parts (iii) and (iv), Microsoft does not appear to be providing any information about who is knowledgeable about Microsoft's "business practices" regarding the models or about the "commercialization and use" of the models. Such individuals would plainly have knowledge relevant to the subject matter of this case, including knowledge relevant to The Times's vicarious copyright infringement claims and to any fair use defenses.

For Interrogatory 7, which seeks information concerning the "existence, custodian, location, and general description of technical documentation and documentation of training processes," Microsoft only provides the release dates of three models: GPT-3.5, GPT-4, and GPT-4 Turbo. R&Os at 11. Microsoft's response is not responsive to this request and fails to address even the foundational issue as to whether such documentation exists. The Times may seek this information through interrogatories under the plain language of SDNY Local Civil Rule 33.3, which permits interrogatories seeking "the existence, custodian, location and general description of relevant documents." Please amend your response to fully answer Interrogatory 7.

For Interrogatories 3, 8, and 9, Microsoft improperly limits its response to *Microsoft's* products, excluding any of OpenAI's products that Microsoft helped create or about which it has relevant knowledge. Interrogatory 3 requests identification of individuals with knowledge about Defendants' practices, policies, abilities, and efforts to track users' interactions with Defendants' Generative AI Products and Services. Microsoft limits its response to individuals with knowledge of "Microsoft's Generative AI Products and Services only." R&Os at 8. But Microsoft does not appear to dispute that individuals employed by Microsoft with knowledge of OpenAI's products and services would have relevant information. Nor could it, as such individuals would have knowledge relevant to, among other topics, The Times's claims of vicarious copyright infringement and to any fair use defense. Instead, Microsoft's objection is that it may not have "comprehensive and accurate information about OpenAI's activities." R&Os at 10. That objection

is not valid and does not absolve Microsoft of its obligation to identify knowledgeable individuals within its organization, which is relevant to the parties' ongoing negotiations regarding search terms and custodians.

Microsoft's response to Interrogatory 8, which requests identification of all of Defendants' Generative AI Products and Services that incorporate, use, or rely on the Text Generation AI Models, is plagued by the same issue. Here, Microsoft again limits its response to identifying only *Microsoft* products that incorporate Text Generation AI Models. R&Os at 12. Does Microsoft have knowledge of which OpenAI products incorporate those models? Please amend your response accordingly, or confirm that Microsoft has no knowledge of which OpenAI products incorporate the models. Additionally, Microsoft's response does not specify "the existence, custodian, location, and general description of documents concerning user engagement, output, user feedback, and analytics" as the interrogatory requests. As above, Local Rule 33.3 permits interrogatories seeking such information. If Microsoft possesses such information, please amend your response accordingly, or confirm in writing that no such documents exist.

Likewise, for Interrogatory 9, which seeks information concerning the existence, custodian, location, and general description of documents reflecting user sessions on Defendants' Generative AI Products, Microsoft states only that "[i]nformation reflecting user sessions on *Microsoft*'s Generative AI Products and Services is maintained within a Search Log API (SLAPI) within Microsoft," and offers to meet and confer regarding the remainder of the interrogatory. R&Os at 13 (emphasis added). The Times takes up that invitation. Please come to the meet and confer prepared to discuss Microsoft's basis for refusing to produce the responsive information and the location of any other documents or information responsive to this interrogatory.

a. Interrogatories for which Microsoft only offers to meet and confer (Interrogatories Nos. 6, 10)

Interrogatory 6 requests Microsoft to identify individuals with knowledge of document storage and preservation as it pertains to the models and products at issue in this case. The scope of this Interrogatory is appropriate, and this information is relevant to understanding the locations and custodians of responsive documents and data. This request is also particularly appropriate considering that the documents produced by Microsoft to date show that the training datasets are stored on Microsoft's servers. As above, Microsoft's objection that it lacks "comprehensive and accurate information about OpenAI's activities" is improper and does not absolve Microsoft of the responsibility of identifying individuals at Microsoft with knowledge of those activities. Please confirm that Microsoft will provide a substantive response to this interrogatory.

Interrogatory 10 requests the custodian, location, and a general description of technical documents related to each training dataset. Regardless of whether Microsoft or OpenAI conducted each phase of training, such information is relevant to The Times's direct and contributory copyright infringement claims, not unduly burdensome to produce, and appropriately sought pursuant to Local Rule 33.3. Please confirm that Microsoft will provide a substantive response to this interrogatory.

**b. Interrogatory for which Microsoft refuses to provide a response
(Interrogatory No. 5)**

Interrogatory 5 requests Microsoft to identify individuals with knowledge of personnel changes within teams overseeing the creation and commercialization of the AI models. This information is relevant to The Times's allegations of willful infringement, Compl. ¶¶ 124–26, particularly considering the numerous, highly publicized reports of terminations and resignations over issues including AI safety and ethics,¹ which are squarely implicated in this case. Microsoft's refusal to provide a response to this interrogatory is improper. Please confirm that Microsoft will provide a response to this interrogatory.

2. Responses and Objections to The Times's Second Requests for Production

a. Global issues

i. Limitation to documents concerning Microsoft's role in training

Microsoft objects to over a dozen requests on the basis that the information sought relates to aspects of the training process completed by OpenAI, and that information is therefore in OpenAI's possession, custody, and control. *See, e.g.*, R&Os to RFPs 16–25, 29, 32–36, 40, 51, 58, and 62. This objection is improper. Any documents that Microsoft possesses relating to any part of the training process—including pre-training, training, mid-training, fine-tuning, and post-training—are relevant, regardless of whether OpenAI or Microsoft primarily oversaw that aspect of training. Microsoft's contention that it “did not Train OpenAI’s models, does not possess the Training Datasets OpenAI used to Train its models, and is unaware of the specific contents of the Training Datasets used by OpenAI to Train its models,” R&Os at 5, does not absolve Microsoft of its obligation to search for and produce all responsive documents that are within its possession, custody or control. Please confirm that Microsoft will search for and produce all such documents.

ii. Limitation to Microsoft's Generative AI Products and Services

Microsoft objects to dozens of requests on the basis that it will “respond [to The Times's requests] with respect to Microsoft’s Generative AI Products and Services only,” or otherwise limits its response to Microsoft's, and not OpenAI's, conduct. *See, e.g.*, R&Os to RFPs 23, 28, 31, 35, 37–43, 45, 46, 48, 50–52, 58–59, 61–67, and 69–71. As the parties have previously discussed, these limitations are improper. The Times's claims concern not only individual products created by Microsoft and OpenAI, but also the technical and financial collaborations between the two companies. Moreover, The Times's contributory infringement claim against Microsoft seeks to hold Microsoft accountable for contributing to and assisting copyright infringement committed by OpenAI. Therefore, any documents that Microsoft has in its possession related to, for example, “measures taken to conceal or obscure Defendants’ use of Times Content during any stage of the development of each of Defendants’ Generative AI Products,” RFP 28, are relevant—regardless of whether that concealment relates to a Microsoft product like Bing Chat or an OpenAI product

¹ Chris Morris, *OpenAI Resignations Are Reaching an Alarming Level. Here Are 11 Key People Who Have Left*, FAST COMPANY (May 17, 2024), <https://www.fastcompany.com/91126785/openai-resignations-are-reaching-an-alarming-level-here-are-11-key-people-who-have-left>.

like ChatGPT. Likewise, any knowledge that Microsoft has regarding OpenAI's use of crawlers to access third-party websites containing Times content, RFP 59, is clearly relevant. Please confirm that Microsoft will search for and produce any responsive documents within Microsoft's possession, including those that relate to OpenAI's conduct and products, including ChatGPT and Browse with Bing.

iii. Limitation to models and datasets “as implemented” in Microsoft’s Generative AI Products and Services (RFPs 23–25, 35, and 40)

Microsoft objects that it will produce documents responsive to The Times's requests regarding datasets and models “as implemented” in Microsoft’s products. R&Os to RFPs 23–25, 35, 40. This objection is improper. Documents regarding, for example, the use of Times content for fine-tuning (RFP 23) and the incremental value of Times content in the models (RFP 24) is relevant regardless of whether those datasets and models were ultimately “implemented” in Microsoft’s products, in OpenAI’s products, or not implemented at all. Please confirm that Microsoft will produce all responsive documents regardless of the implementation of certain datasets and models in Microsoft’s products.

iv. Limited date range (RFPs 35 and 48)

In response to RFP 35, which asks about changes made to Defendants’ models and products from September 2021 to present, and RFP 48, which asks about user and subscription data during the same period, Microsoft limits its cut-off date to December 28, 2023, the day after The Times filed its complaint. This cut-off is arbitrary and improper. Please confirm that Microsoft will amend its response and continue to supplement its document productions accordingly throughout the discovery period for this case.

v. Objections that RFPs are duplicative of Requests for Inspection (RFPs 19–22, 38, 42, 45, 47)

Microsoft objects that certain Times Requests for Production are duplicative of its Requests for Inspection. R&Os to RFPs 19–22, 38, 42, 45, 47. This is improper because both categories of requests seek different information and serve different purposes. The Requests for Inspection permit The Times to inspect data, source code, models, and products, whereas the Requests for Production seek information and documents in Defendants’ possession, custody, and control that Microsoft bears the burden to search for and produce. While related data for each request may be obtained through The Times’s Requests for Inspection, Microsoft cannot shift its own discovery obligations to The Times. Please confirm that Microsoft will produce all responsive documents, regardless of whether the subject matter of those documents relates to The Times’s Requests for Inspection.

vi. Improper use of “sufficient to show” in discovery responses (RFPs 24, 25, 39)

On several occasions, Microsoft improperly limits its production to documents “sufficient to show” the requested information. This is inadequate. For example, RFPs 24 and 25 request

documents concerning the incremental value and analyses of the placement of journalistic content, including Times content, in training datasets. Microsoft responds that it will produce documents “sufficient to show” the requested information as it pertains to “Microsoft’s Fine-Tuning Datasets as implemented in Microsoft’s Generative AI Products and Services.” In addition to suffering from the other global deficiencies laid out above, Microsoft’s responses to these requests improperly limit Microsoft’s production to a mere subset of relevant documents of Microsoft’s choosing. Any self-imposed “sufficient-to-show” limitation is improper because that approach “would let [Microsoft] decide what the answer[s] [are] on the merits [of factual issues in the case] and then limit discovery accordingly.” *In re Stubhub Refund Litig.*, 2022 WL 1640304, at *2 (N.D. Cal. May 24, 2022). Please confirm that Microsoft will produce all relevant documents responsive to these requests.

Further, RFP 39 requests documents concerning testing for hallucinations and memorization. Here too, Microsoft agrees to produce only a subset of documents “sufficient to show” the requested information. Again, this is improper because it allows Microsoft to exclude highly relevant and responsive information, such as internal communications about the frequency and severity of these known issues. And further, Microsoft’s response omits reference to sub-parts (ii)–(v). Please confirm that Microsoft did not intend this omission to be limiting and that it will produce all documents responsive to this RFP, and please revise your response to address sub-parts (ii)–(v).

b. Issues with Microsoft’s responses to specific RFPs

i. Requests for which Microsoft refuses to produce documents (RFPs Nos. 29, 31, 46, 56, 58, 60, 62, 63, and 67)

Microsoft has refused to produce documents in response to a number of The Times’s requests. *See R&Os to RFPs 29, 31, 46, 56, 58, 60, 62, 63 and 67.* This is improper, and The Times requests to meet and confer regarding Microsoft’s refusal to respond to these relevant inquiries.

RFP 29 requests documents concerning Defendants’ reliance on earlier versions of Defendants’ Generative AI Models to be used in or to train or develop later versions of Defendants’ Generative AI Models. Microsoft’s objection that this request could “encompass multiple processes” is without merit. To the extent that, at any phase of the pre-training, training, mid-training, fine-tuning, or post-training process, the models have been built on outputs, weights, and other artifacts of earlier models that themselves were built using copyrighted content, that information is squarely within the scope of this case.

RFP 31 requests documents concerning Defendants’ technical documentation for their Generative AI Products and Services. Microsoft both states that it “will respond with respect to Microsoft’s Generative AI Products and Services only” but also that it “will not produce documents in response to this Request.” The former objection is improper, as discussed above in §2(a)(ii). Microsoft also objects on the bases that the requested information is irrelevant to the claims and defenses in this case and that the request is overly broad and not proportional to the needs of the case. This is incorrect. The technical documentation requested is crucial to The Times’s claims. Moreover, the fact that the definition of “technical documentation” is derived from

international law does not affect its relevance, and indeed, Microsoft should already have these documents readily available to identify and produce based on its obligations under international law.

RFP 46 requests documents concerning how Defendants' products generate follow-up queries for users. Microsoft objects that this request is irrelevant to the claims and defenses in this case, but such information is relevant to Defendants' efforts and abilities to retain user traffic to grow their generative AI business at the expense of copyright holders, which goes to Defendants' fair use defenses.

RFP 56 requests documents concerning terminations and resignations of certain members of relevant teams overseeing issues such as AI safety and ethics. Here, too, Microsoft objects to relevance, but the requested information is relevant to willfulness, among other issues, given the numerous, highly publicized departures during which departing employees have cited internal issues related to, for example, Defendants' treatment of publishers' intellectual property. It is unclear how this information—including information already disclosed by third-parties—might be "confidential and private," R&Os at 43, and Microsoft has not explained how these documents might be privileged. Moreover, confidentiality does not constitute a basis for withholding. Microsoft can produce sensitive information on a confidential basis, as it has done for other requests.

RFP 58 requests documents related to the use of Defendants' Generative AI Products and Services to promulgate misinformation and disinformation. Microsoft objects to The Times's use of the term "pink-slime" journalism as vague. For the avoidance of doubt, this phrase refers to partisan fake news sites that are often powered by Defendants' APIs and other Generative AI Products and Services.² These topics are relevant to contributory infringement, fair use, and tarnishment of The Times's trademarks, among other issues.

RFP 60 requests documents concerning indemnification, common interest agreements, and joint defense agreements. Microsoft objects that the request is overbroad because it seeks documents concerning common interest agreements, joint defense agreements, or indemnification obligations and that the requested documents are not relevant to the claims and defenses in this Action. Both assertions are incorrect, as The Times's request specifically seeks agreements "between Defendants related to Defendants' Generative AI Products or Services," and the requested documents are relevant to The Times's vicarious and contributory infringement allegations, among other issues. Microsoft also objects that the requested information is privileged. While some requested documents may be privileged, that does not relieve Microsoft of its obligation to search for and produce all responsive, non-privileged documents and log withheld documents.

RFPs 62 and 63 request documents concerning Defendants' competition and communications with other LLM creators regarding the use of copyrighted content in its Generative AI Models, Products, and Services. Contrary to Microsoft's objections, these requests

² Jack Brewster, *How I Built an AI-Powered, Self-Running Propaganda Machine for \$105*, WALL STREET JOURNAL (Apr. 12, 2024), <https://www.wsj.com/politics/how-i-built-an-ai-powered-self-running-propaganda-machine-for-105-e9888705>.

are targeted to documents that are relevant to The Times's willfulness allegations, Microsoft's knowledge of infringing use, Defendants' motives for using copyrighted content, and the benefits Defendants derived from using The Times's copyrighted content.

RFP 67 requests documents related to Defendants' business plans, both for: (i) Defendants' Generative AI Products and Services; and (ii) future products and services in the categories of news, cooking, consumer reviews, and sports. Such information is relevant to fair use and the substitute effect of Defendants' products on journalism, among other issues.

Please let us know whether Microsoft is willing to reconsider its refusal to produce any documents responsive to these requests, all of which seek information relevant to and squarely within the scope of this case.

ii. Requests for which Microsoft is willing to meet and confer (RFPs Nos. 19, 21, 27, 32, 33, 37, 38, 41, 42, 45, 47, 49, 51, 53, 64, and 68)

There are a number of requests for which Microsoft has currently not agreed to produce the requested documents but instead has offered to meet and confer.

Microsoft's objections to RFPs 19, 21, 32, and 33—that these requests seek information related to training that OpenAI conducted—are improper for the reasons stated above. Microsoft's objections to RFPs 19, 38, 42, 45, and 47—that these requests seek information related to The Times's requests for inspection—are likewise improper for the reasons stated above.

RFP 27 requests documents concerning the use of robots.txt, and Microsoft objects that this request is overbroad and irrelevant. The Times disagrees. Specifically, The Times seeks information regarding Microsoft's policies and practices related to third-party robots.txt directives and/or metatags that block the crawling and scraping of copyrighted content and whether Defendants rely on extending robots.txt permissions surrounding crawling and indexing for Search for GAI purposes, issues that are relevant to the extent that the models are trained to use or avoid using copyrighted content. The extent to which Bing respects crawler directives directly impacts the data available to GPT models and the content available for model responses, which are clearly implicated in this case.

RFP 37 requests core technical documents related to Defendants' generative AI models—information that is again squarely within the scope of this case. Microsoft objects that this request is overbroad because it seeks technical documents “regarding any material, at any time, and without any subject matter limitation.” This is incorrect. The Times's request specifically asks for this information with respect to Defendants' Generative AI Products and Services.

RFPs 41, 42, 45, 47, and 51 request documents related to the interactions between users and Defendants' generative AI products. These requests include studies related to users' trust in outputs (RFP 41), which relates to the benefits of using high-quality content such as Times content for training and fine-tuning; Defendants' efforts and abilities to monitor the use of Defendants' products (RFPs 42, 45), which relates to Defendants' knowledge of copyrighted outputs and efforts by users to circumvent access control measures using Defendants' products; and user engagement

metrics (RFP 47, 51), which relate to Defendants' efforts and abilities to retain user traffic and grow their generative AI business.

RFP 49 seeks documents concerning Defendants' monitoring of The Times's investigation into whether Defendants' generative AI models contain copyrighted content. Microsoft responds that it would like to meet and confer, especially with respect to the meaning of "contain." For the avoidance of doubt, "contain" refers to whether Times content is encoded, memorized, stored, or otherwise fixed in the model such that it can be reproduced from the model.

RFP 53 requests documents related to the decision to exclude Bing Chat referral traffic data from Bing Webmaster Tools.³ Such documents are relevant because the exclusion of Bing Chat referral traffic from Bing Webmaster Tools relates to the extent to which Bing Chat—a product indisputably at issue in this case—is referring traffic to websites directly and Microsoft's decision to exclude Bing Chat referral traffic from Bing Webmaster Tools relates to Microsoft's knowledge about such referral traffic. Search referral information, and the decision to stop providing that information, is relevant to, among other issues, the substitutive effect of Defendants' products on journalism.

The Times accepts Microsoft's invitation to discuss these issues further on a meet and confer and requests that Microsoft confirm its agreement to produce any relevant, responsive documents.

iii. Requests for which Microsoft's responses are incomplete or require clarification (RFPs Nos. 16, 18, 22, 35, 40, 52, 54, 55, 61 66, 69, and 70)

Microsoft has provided incomplete responses to a number of The Times's requests, making it unclear to what extent Microsoft intends to search for and provide relevant, responsive documents. This is improper, and The Times requests a meet and confer regarding these issues.

RFP 16 requests documents concerning the decision not to publicly release certain information regarding models that were created after GPT-3. Microsoft has agreed to search for and produce documents "concerning communications between Microsoft and OpenAI regarding the public disclosure of the contents of OpenAI's Training Datasets for its Text Generation AI Models. It is unclear from Microsoft's response whether it is agreeing to search for internal Microsoft communications on this topic. Please confirm that Microsoft will also search for and produce any internal communications that are responsive to this request.

RFP 18 requests documents related to the removal of copyright management information. Microsoft responds that it will produce documents related to the *post-training* removal of copyright management information. This is improper, as the removal of copyright management information at any stage of the training process is relevant to this lawsuit. Please confirm that Microsoft will

³ *Building the New Bing*, MICROSOFT BING BLOGS (Feb. 21, 2023), <https://blogs.bing.com/search-quality-insights/february-2023/Building-the-New-Bing>; Barry Schwartz, *Bing Chat Data Still Not In Bing Webmaster Tools*, SEARCH ENGINE ROUNDTABLE (Aug. 23, 2023), <https://www.seroundtable.com/bing-webmaster-tools-bing-chat-data-missing-35928.html>.

search for and produce relevant documents related to the removal of copyright management information at any stage of the training process.

RFP 22 requests documents sufficient to identify each Times article contained within Defendants' training datasets. Microsoft responds that it will produce "documents sufficient to show Times articles contained within Microsoft's fine-tuning datasets." This response is insufficient in two ways. First, The Times seeks documents sufficient to show articles contained within all datasets, not just Microsoft's fine-tuning datasets. Please confirm that Microsoft will produce documents sufficient to identify all Times articles contained in all datasets, or confirm that Microsoft is not in possession, custody, or control of any datasets other than fine-tuning datasets. Second, The Times seeks documents sufficient to show *each* Times article contained therein and the frequency with which they appear. Please confirm that the documents Microsoft has agreed to produce will reflect the frequency in which each article of Times Content appears in its training datasets.

RFP 35 requests documents concerning any changes made to Defendants' generative AI models and products. As discussed, Microsoft's response is insufficient in its limitation to documents related to fine-tuning. Moreover, please be prepared to explain during the meet and confer what is meant by documents "relating to the implementation of additional mitigations to prevent the output of copyrighted material and/or improper attributions," and what limitation, if any, this language is intended to convey.

RFP 40 requests documents related to the prevalence of outputs that contain copies of text included in the Training Datasets, ***including*** subparts (i), (ii), and (iii). Microsoft's response states that it will search for and produce documents responsive to the three subparts but does not agree to otherwise search for and produce documents concerning the prevalence of outputs that contain copies of training data. Please confirm that Microsoft will conduct a search for documents related to the prevalence of outputs that contain copies of text included in the Training Datasets.

RFP 52 requests documents related to the appearance of Times content in generative AI outputs. In response, Microsoft states that it will produce what it agreed to produce in response to RFP 45, R&Os at 41, but RFP 45 relates to user queries. Please confirm that Microsoft intended to reference RFP 43 ("Any surveys, studies, analyses, or statistics on patterns and practices any of Defendants' Generative AI Products and Services, including any mentions of The New York Times, such surveys, studies, analyses, or statistics regarding of user engagement with in user queries and/or inclusion of links to Times Content in user queries or outputs") in its response, and please revise the response accordingly.

RFP 54 requests documents related to the decision to suspend Browse with Bing. Microsoft responds that it "will conduct a reasonable search and produce responsive, non-privileged documents, if any, OpenAI's decision to disable Browse with Bing in July 2023." R&Os at 42. Please confirm that Microsoft intended to state that it would produce documents "*relating to*" this decision.

RFP 55 seeks documents concerning any public statements made by Microsoft or its employees regarding this litigation. Microsoft states that it will produce documents concerning

such statements made by “Microsoft itself.” R&Os. Please be prepared to explain during the meet and confer what constitutes a statement made by “Microsoft itself,” and in what ways, if any, this response is intended to be limiting.

RFP 61 requests documents related to negotiations between Defendants and content owners over the use of copyrighted content in Defendants’ models, products, and services. In response, Microsoft has only agreed to produce documents related to negotiations with The Times, despite the clear relevance of information regarding negotiations with other content owners to The Times’s damages calculations, fair use, and other issues. Please confirm that Microsoft intends to produce documents concerning licensing and negotiations with all content owners, not solely The Times.

Finally, Microsoft’s responses to several RFPs calling for financial documents suggest that Microsoft may be excluding relevant, responsive documents from its productions, as follows:

- Sub-part (iii) of RFP 66 requests “product-level revenue data from 2021 to present.” Microsoft’s response does not address this portion of The Times’s request;
- Similarly, Microsoft’s response to RFP 69 does not address The Times’s request for “referral” and “affiliate” revenue; and
- Microsoft’s response to RFP 70 does not address The Times’s requests related to costs and profit margins, despite the clear relevance of these requests to The Times’s damages computations.

Please confirm that Microsoft will produce the requested financial information, including at the product-level where available.

The Times appreciates Microsoft’s attention to these issues and requests a meet and confer no later than July 26, 2024.

Respectfully,

/s/ Zach Savage
Zach Savage